

United States District Court,
E.D. Texas, Marshall Division.

OPTI INC,
Plaintiff.

v.

NVIDIA CORPORATION,
Defendant.

No. Civ.A.2:04-CV-377TJW

April 24, 2006.

Samuel Franklin Baxter, Attorney at Law, Marshall, TX, Danielle A.R. Coffman, David Thomas Bower, Michael L. Brody, Taras Alexander Gracey, Winston & Strawn, Chicago, IL, David J. Cerveny, Steven M. Bauer, Proskauer Rose LLP, Boston, MA, Gary Scott Kitchen, Martin C. Robson, III, McKool Smith, Dallas, TX, Kevin Lee Burgess, McKool Smith, Austin, TX, for Plaintiff.

Iain R. Cunningham, Cooley Godward, San Francisco, CA, Jessica Valenzuela Santamaria, Lowell D. Mead, Matthew J. Brigham, Sarah J. Guske, Stephen C. Neal, Timothy S. Teter, Cooley Godward, Palo Alto, CA, John Michael Pickett, Damon Michael Young, Lance Lee, Young Pickett & Lee, Texarkana, TX, Michael G. Rhodes, Cooley Godward, San Diego, CA, for Defendant.

MEMORANDUM OPINION AND ORDER

WARD, J.

I. Introduction

Plaintiff OPTi Inc. ("OPTi") has asserted that defendant nVidia Corporation ("nVidia") infringes two families of OPTi patents:

-> OPTi's "Compact ISA" or "CISA" patents-Namely, claims 1, 12, 13, 19, and 21 of United States Patent No. 6,098,141 ("the '141 patent") and claim 5 of United States Patent No. 5,944,807 ("the '807 patent"); and

-> OPTi's "Pre-Snoop" patents-Namely, claims 1, 7, 8, 9, 21, and 26 of United States Patent No. 5,710,906 ("the '906 patent"); claims 1, 3, 9, 17, 18, and 26 of United States Patent No. 5,813,036 ("the '036 patent"); and claims 1, 9, 11, 33, 40, 42, 44, and 46 of United States Patent No. 6,405,291 ("the '291 patent").

The Court issues this Order to resolve the parties' claim construction disputes for the asserted patents.

II. Legal Principles of Claim Construction

"A claim in a patent provides the metes and bounds of the right which the patent confers on the patentee to exclude others from making, using or selling the protected invention." *Burke, Inc. v. Bruno Indep. Living Aids, Inc.*, 183 F.3d 1334, 1340 (Fed.Cir.1999). Claim construction is an issue of law for the court to decide. *Markman v. Westview Instruments, Inc.*, 52 F.3d 967, 970-71 (Fed.Cir.1995) (en banc), *aff'd*, 517 U.S. 370, 116 S.Ct. 1384, 134 L.Ed.2d 577 (1996).

To ascertain the meaning of claims, the court looks to three primary sources: the claims, the specification, and the prosecution history. *Markman*, 52 F.3d at 979. Under the patent law, the specification must contain a written description of the invention that enables one of ordinary skill in the art to make and use the invention. A patent's claims must be read in view of the specification, of which they are a part. *Id.* For claim construction purposes, the description may act as a sort of dictionary, which explains the invention and may define terms used in the claims. *Id.* "One purpose for examining the specification is to determine if the patentee has limited the scope of the claims." *Watts v. XL Sys., Inc.*, 232 F.3d 877, 882 (Fed.Cir.2000).

Nonetheless, it is the function of the claims, not the specification, to set forth the limits of the patentee's claims. Otherwise, there would be no need for claims. *SRI Int'l v. Matsushita Elec. Corp.*, 775 F.2d 1107, 1121 (Fed.Cir.1985) (en banc). The patentee is free to be his own lexicographer, but any special definition given to a word must be clearly set forth in the specification. *Intellicall, Inc. v. Phonometrics*, 952 F.2d 1384, 1388 (Fed.Cir.1992). And, although the specification may indicate that certain embodiments are preferred, particular embodiments appearing in the specification will not be read into the claims when the claim language is broader than the embodiments. *Electro Med. Sys., S.A. v. Cooper Life Scis., Inc.*, 34 F.3d 1048, 1054 (Fed.Cir.1994).

This court's claim construction decision must be informed by the Federal Circuit's decision in *Phillips v. AWH Corp.*, 415 F.3d 1303 (Fed.Cir.2005)(en banc). In *Phillips*, the court set forth several guideposts that courts should follow when construing claims. In particular, the court reiterated that "the *claims* of a patent define the invention to which the patentee is entitled the right to exclude." *Id.* at 1312 (emphasis added)(quoting *Innova/Pure Water, Inc. v. Safari Water Filtration Sys., Inc.*, 381 F.3d 1111, 1115 (Fed.Cir.2004)). To that end, the words used in a claim are generally given their ordinary and customary meaning. *Id.* The ordinary and customary meaning of a claim term "is the meaning that the term would have to a person of ordinary skill in the art in question at the time of the invention, i.e. as of the effective filing date of the patent application." *Id.* at 1313. This principle of patent law flows naturally from the recognition that inventors are usually persons who are skilled in the field of the invention. The patent is addressed to and intended to be read by others skilled in the particular art. *Id.*

The primacy of claim terms notwithstanding, *Phillips* made clear that "the person of ordinary skill in the art is deemed to read the claim term not only in the context of the particular claim in which the disputed term appears, but in the context of the entire patent, including the specification." *Id.* Although the claims themselves may provide guidance as to the meaning of particular terms, those terms are part of "a fully integrated written instrument." *Id.* at 1315 (quoting *Markman*, 52 F.3d at 978). Thus, the *Phillips* court emphasized the specification as being the primary basis for construing the claims. *Id.* at 1314-17. As the Supreme Court stated long ago, "in case of doubt or ambiguity it is proper in all cases to refer back to the descriptive portions of the specification to aid in solving the doubt or in ascertaining the true intent and meaning of the language employed in the claims." *Bates v. Coe*, 98 U.S. 31, 38, 25 L.Ed. 68 (1878). In addressing the role of the specification, the *Phillips* court quoted with approval its earlier observations from *Renishaw PLC v. Marposs Societa' per Azioni*, 158 F.3d 1243, 1250 (Fed.Cir.1998):

Ultimately, the interpretation to be given a term can only be determined and confirmed with a full understanding of what the inventors actually invented and intended to envelop with the claim. The construction that stays true to the claim language and most naturally aligns with the patent's description of the invention will be, in the end, the correct construction.

Consequently, *Phillips* emphasized the important role the specification plays in the claim construction process.

The prosecution history also continues to play an important role in claim interpretation. The prosecution history helps to demonstrate how the inventor and the PTO understood the patent. *Phillips*, 415 F.3d at

1317. Because the file history, however, "represents an ongoing negotiation between the PTO and the applicant," it may lack the clarity of the specification and thus be less useful in claim construction proceedings. *Id.* Nevertheless, the prosecution history is intrinsic evidence. That evidence is relevant to the determination of how the inventor understood the invention and whether the inventor limited the invention during prosecution by narrowing the scope of the claims.

Phillips rejected any claim construction approach that sacrificed the intrinsic record in favor of extrinsic evidence, such as dictionary definitions or expert testimony. The *en banc* court condemned the suggestion made by *Tex. Digital Sys., Inc. v. Telegenix, Inc.*, 308 F.3d 1193 (Fed.Cir.2002), that a court should discern the ordinary meaning of the claim terms (through dictionaries or otherwise) before resorting to the specification for certain limited purposes. *Id.* at 1319-24. The approach suggested by *Tex. Digital*-the assignment of a limited role to the specification-was rejected as inconsistent with decisions holding the specification to be the best guide to the meaning of a disputed term. *Id.* at 1320-21. According to *Phillips*, reliance on dictionary definitions at the expense of the specification had the effect of "focus[ing] the inquiry on the abstract meaning of words rather than on the meaning of the claim terms within the context of the patent." *Id.* at 1321. *Phillips* emphasized that the patent system is based on the proposition that the claims cover only the invented subject matter. *Id.* What is described in the claims flows from the statutory requirement imposed on the patentee to describe and particularly claim what he or she has invented. *Id.* The definitions found in dictionaries, however, often flow from the editors' objective of assembling all of the possible definitions for a word. *Id.* at 1321-22.

Phillips does not preclude all uses of dictionaries in claim construction proceedings. Instead, the court assigned dictionaries a role subordinate to the intrinsic record. In doing so, the court emphasized that claim construction issues are not resolved by any magic formula. The court did not impose any particular sequence of steps for a court to follow when it considers disputed claim language. *Id.* at 1323-25. Rather, *Phillips* held that a court must attach the appropriate weight to the intrinsic sources offered in support of a proposed claim construction, bearing in mind the general rule that the claims measure the scope of the patent grant. The court now turns to a discussion of the claim construction disputes.

II. Description of the Technology

The OPTi patents relate to "core logic" chipsets, the processors that direct traffic between the central processor, memory, input/output devices, graphics cards, video cards, and various other devices that are contained within, or connected to, a computer. In particular, the CISA and Pre-Snoop patents relate to inventions that make efficient use of two valuable commodities in data processing: space and time.

1. Background

In the earliest days of computer processing, there were no core logic chipsets. The central processor communicated directly with peripheral devices that made up the computer. As computers got more complicated, chipsets were introduced as a way of coordinating the burgeoning array of functionality and relieving central processors of that administrative burden. This freed more CPU resources for the fundamental mission of computing.

Broadly speaking, a typical chipset operates as an input/output (I/O) hub for the CPU, memory, peripherals, etc. The typical chipset consists of two processors: a Northbridge and a Southbridge. The Northbridge communicates with the CPU, the system memory and the graphics card. The Southbridge communicates with the peripheral devices attached to, and incorporated in, the computer. The Northbridge and the Southbridge communicate with one another.

The links between the various devices comprising the computer, known as interfaces, consist of conductors on which the devices transmit signals to one another, communicating address, command, and data

information. The most common type of interface is known as a bus. The buses and interfaces allow the various computer devices to exchange data and to operate in coordination with one another.

The inventions of the CISA and Pre-Snoop patents arose from OPTi's efforts to maximize the efficiency of computers using its chipsets that transferred data over two types of buses: the Industry Standard Architecture ("ISA") bus, and the Peripheral Component Interconnect ("PCI") bus. The ISA bus emerged in the early 1980s with IBM's AT line of computers that operated with an Intel 80286 processor. The PCI bus, which first appeared a decade later in computers using Intel's 80486 processors, was intended to replace the ISA bus.

2. The Compact ISA Patents

Though the PCI bus was intended to replace the ISA bus, the need to support legacy peripherals and new low cost devices gave the ISA bus an unanticipated longevity. As originally designed, the ISA bus dedicated separate conductors to transmit signals for each of the functions that the devices attached to it wanted to perform. Thus, when a device connected to the ISA bus wanted to transfer data to or from another device in the computer, it would use a dedicated "address" line to transmit the address of the "target" device with which it wished to communicate, a separate line dedicated to transmitting a "command" signal identifying the type of transaction it wanted to engage in, and still other lines dedicated for use in transmitting the data itself.

The OPTi patents were generally focused on addressing the need of having signal lines available to handle the ISA devices, the need to have signal lines to handle the requirements of PCI devices, and at the same time, minimizing the amount of space taken up by this complex of interconnections. OPTi used "multiplexing" to solve this problem.

Multiplexing is a concept used in many industries that involves the transmission of signals from one place to another. Instead of dedicating a single transmission channel to a single signal governing a single function, the channel will transmit multiple signals used for multiple purposes at different times, different frequencies, or different wavelengths, etc.

For example, where the channel is an electrical circuit, the different devices using the circuit may be assigned time slots so that the line is used for one purpose for a fixed period, then for another purpose, then a third, then the first one picks up where it left off and so on. This is known as time division multiplexing. Various forms of multiplexing are available depending on the medium of transfer, the devices and the functions involved. By using multiplexing, OPTi was able to reduce the number of signal lines needed on an ISA bus from 58 to 22, and correlatively, to reduce the size of its chipsets, by reducing the number of "pins" attached to the reduced number of signal lines.

3. The Pre-Snoop Patents

The Pre-Snoop patents addressed a second issue that arose with the introduction of the PCI bus and the subsequent development of the Pentium and Pentium-compatible processors. One of the advantages of the PCI was its ability to transfer data from one device to another by a particular method called "burst" transfers. The Pre-Snoop patents disclose a technique for optimizing such burst transfers with Pentium processors.

Data is stored, created or used at a lot of places in a computer. Each such location is known as an address. For example, a memory storage device containing data to be read (the "target") cannot know that it is being asked to transfer data or what data to transfer unless and until the requesting device (the "master") puts an address onto the bus notifying the target that it is the object of a request and notifying the target what data is being requested. In a "burst" transfer, this information is all that the target needs to figure out which data to

transfer, as the target dispatches data until the target is told to stop by the master or elects to stop the transaction itself.

A complication arises in this scenario because much of the memory can be stored in two places: main memory or cache memory. Cache memory is memory that stores copies of information expected to be used by the CPU at addresses that correspond to addresses for that information in secondary memory. This memory typically operates at particularly high speed and is typically positioned adjacent to the CPU. Access to the cache is thus generally quicker than access to the main memory. This speeds up the CPU's ability to access and process the data that it needs.

As the CPU processes data, it saves that data to the cache for continued convenient access. The problem is that the CPU may well change the data that it is processing. If that modified data is stored only in the cache, it will not be identical to the data stored on, for example, the disk drive from which it was initially read. Thus, if some other device—a CD drive, for example—accesses the main memory to read data, it may get data that is no longer current.

In Intel's X86 line of CPU's, the system solved this problem by using a "write-through" cache. Basically, as data was modified by the CPU, it was written to both the main and cache memories, thereby assuring constant "cache consistency." In the Pentium processors, the cache was a "write-back" cache. This meant that the CPU did not take the time to write every modification through to main memory. Instead, modified data was stored in the cache with a flag to indicate its modified state. Thus, to read data from memory in a Pentium system, it was necessary to adopt some mechanism to check for the presence of this flag, to assure that the data in main memory was valid, and that the cache did not contain a version of data that had been modified by the CPU.

The mechanism adopted was a "snoop." For example, a bus master seeking to initiate a transaction would first initiate an "inquire" or "snoop" cycle to the CPU to find out whether the data being sought had been stored in cache memory in a modified state and to "write back" the most current version of the data to the main memory if a modification had been made.

The PCI standard required one "line" of cache memory to be snooped at a time, and then permitted transfer of that line if the snoop showed it to be either absent from the cache or in the cache but unmodified. The PCI protocol required that a transfer stop at the end of each line transferred, snoop the next line, transfer the line just snooped, and then stop again to snoop the succeeding line. Because burst transfers could encompass any amount of data, including data stored in multiple lines of memory, this practice resulted in a non-uniform transfer in which the bus spent more time sitting idle than it did carrying data.

The sole exception to this rule was that multiple lines of memory could be read without stopping when the data to be transferred was not cacheable. In this instance, the snoop operation could be ignored because there was no risk of stale data being accessed. Thus, the entire burst of data could be sequentially pre-fetched and read to the master without interruption.

The Pre-Snoop patents embody the idea that cacheable memory could be transferred nearly as rapidly as non-cacheable memory if the snoop of a line were conducted while the preceding line was transferring. In that way, as in the case of non-cacheable memory, the system would know that the line was not stale and there would be no need to stop the transfer at the end of the first line to snoop the second line because that snoop would already be completed. The "snoop ahead" process could be repeated as long as the burst transfer was underway so that all of the data comprising the burst could be sent without interruption and at a constant rate.

III. Analysis of the Disputed Terms


A. CISA Patents

1. host platform

OPTi contends that "host platform" is "a bus controller to which a CPU is connected by a CPU or host bus and to which peripheral devices are connected by a peripheral bus." nVidia contends that "host platform" is "a device or devices which include one or more CPUs, to which other devices (peripherals) are connected and that generally controls those devices." While using somewhat different language, the parties' proposed constructions are similar. The principal difference is whether the host platform is connected to a CPU (OPTi's construction) or whether a host platform must include a CPU (nVidia's construction).

To support its construction, nVidia first notes that in prior litigation, FN1 OPTi recognized that the "host platform" included one or more CPUs. In addition, nVidia turns to Figure 1 of the '141 patent (shown below).

FN1. *OPTi v. National Semiconductor*



Based largely on Figure 1, nVidia asserts that the host platform interface circuit (115) is not the host platform. Instead, the host platform interface circuit forms the interface between the host platform and the peripheral devices. Thus, nVidia contends, the host platform interface circuit in Figure 1 is not the entire host system but is instead only a part of the host platform.

According to nVidia, the host system is the host side (105) in Figure 1 of the patent. *See also* '141 patent, Col. 4:59-62 ("FIG. 1 shows a host side 105 and a device side 110 of an interface 175. On the host side 105 is a host platform interface circuitry 115 coupled to CPU 120 via a CPU bus 125.").

OPTi responds by first noting that the statements by OPTi in prior litigation were made by prior OPTi counsel and that OPTi believes its prior position to have been in error. Consequently, OPTi advances what it believes to be the correct construction now. Further, OPTi notes that National Semiconductor and OPTi resolved their litigation before the Court in that case issued its claim construction opinion. Accordingly, there is no judicial estoppel. In this Court's view, claim construction is a matter of law and the Court will give OPTi's proposed construction the same consideration it provides to nVidia's proposed construction.

OPTi asserts that the intrinsic record supports its construction and also points to Figure 1 and citations in the '141 patent to support its position. OPTi contends that the CPU 120 is coupled to the host platform interface circuit and that, therefore, the CPU and host platform must be distinct from one another.

To further support its argument, OPTi also relies on the problem the CISA patents sought to solve. The origin of the CISA patents was to design a new chipset interface so that the increasing demands of new bus architecture could be managed. Specifically, "[t]he present invention, roughly described, is directed to an interface to be used between a host device and one or more peripheral devices." '141 patent, Col. 3:24-26.

The Court is persuaded that OPTi's construction is correct. In addition to the arguments raised by OPTi, the Court notes that the specification uses "host" and "host platform interface circuitry 115" interchangeably:

Host platform interface circuitry 115 asserts CMD# synchronously with the rising edge of ATCLK. Host can also optionally inhibit its ISA MRD# /MWR # lines.

'141 patent, Col. 7:55-58 (emphasis added); *see also* Col. 9:52-54; Col. 10:1-5. Accordingly, the Court adopts OPTi's proposed construction of "host platform."

2. address-data bus

OPTi's construction of "address-data bus" is "a bus that carries both address information and data." nVidia's construction is similar. nVidia contends that "address-data bus" is "the bus that carries address information in one or more address phases, and data information in one or more data phases." As the parties' proposed constructions are similar, the Court's inquiry for this dispute is focused on whether or not "address-data bus" is limited to the use of phasing techniques.

According to nVidia, the intrinsic evidence defines the address-data bus as sending address information and data in phases. The '141 patent Abstract explains:

Address, data, command, interrupt request, and DMA request information are communicated between the host and the peripheral device via a single bus by multiplexing the information on the bus using phase techniques.

Furthermore, nVidia relies on this Court's precedent that statements in the Summary of the Invention and prosecution history that define the invention may limit the scope of the claims. *See IAP Intermodal, L.L.C. v. Northwest Airlines Corp.*, 2:04-CV-65 (Memorandum Opinion and Order of Sept. 7, 2005). Accordingly, nVidia contends that the Summary of the Invention in the '141 patent is so limiting:

The present invention, roughly described, is directed to an interface to be used between a host device and one or more peripheral devices ... Using phasing techniques, the CISA interface multiplexes the different types of information onto the single bus. For instance, in one cycle, address information may be driven onto the bus for a time period, an "address phase," followed by data information in a "data phase."

'141 patent, Col. 3:24-58. In further support of its argument, nVidia notes that every other example in the

'141 patent specification that communicates address and data information communicates address and data information over the "address-data bus" in at least one address phase and at least one data phase. *See* '141 patent, Col. 6:7-18.

nVidia's proposed construction is incorrect. nVidia is attempting to incorporate limitations from the preferred embodiment into the claims. The citations nVidia relies on to support its construction are full of non-limiting words such as: "[t]he present invention, *roughly described*" and "[f]or instance." These non-limiting words appear throughout the Summary of the Invention and specification.

There is nothing in the Summary of Invention or the rest of the specification that would require "address-data bus" to be limited in the manner nVidia desires. To do so would be error. *See* RF Del., Inc. v. Keystone Techs., Inc., 326 F.3d 1255, 1264 (Fed.Cir.2003). Phasing techniques are just one method of multiplexing information and the claims will not be limited to phasing techniques absent a disavowal of claim scope by the patentee.

In addition, claim differentiation supports OPTi's proposed construction. As a matter of claim construction, dependent claims are generally narrower in scope than the claims from which they depend. *Lampi Corp. v. Am. Power Prods., Inc.*, 228 F.3d 1365, 1376 (Fed.Cir.2000). The doctrine of claim differentiation "is clearly applicable when there is a dispute over whether a limitation found in a dependent claim should be read into an independent claim, and that limitation is the only meaningful difference between the two claims." *Werner Mfg. v. Coating Mach. Sys., Inc.*, 239 F.3d 1225, 1233 (Fed.Cir.2001).

Independent claim 1 of the '141 patent states that the invention comprises "an address-data bus to carry in a multiplexed manner address information ... [and] data information." Dependent claim 2 states that the "said address information is carried on said address-data bus during an address phase ... and said data information is carried during a data phase[.]" Thus, if phasing techniques were a limitation of address-data bus in claim 1, it would render claim 2 of the '141 superfluous.

nVidia asserts that claim differentiation does not apply in this case because claim 2 additionally contains the phrase "said data phase follows said address phase[.]" Thus, nVidia contends, the difference between claim 1 and claim 2 is that the data phase may precede the address phase in claim 1. The Court rejects this argument as it is understood that the data phase follows the address phase. *See* September 4, 1998 Office Action at 3 ("one of ordinary skill in the art would have considered having the data phase follow the address/command phase because following this sequence is important in order to avoid sending data information to the wrong destination/address."). Accordingly, because the patentee did not disavow other methods of multiplexing and claim differentiation undermines nVidia's assertions, the Court adopts OPTi's proposed construction for "address-data bus."

3. address information

OPTi contends that "address information" means "at least a portion of an address." nVidia contends that the correct construction of "address information" is "the bits that identify the intended destination or receiver of a transmission." Thus, the principal debate between the parties on this construction is whether "address information" requires an entire address or just a part of an address.

To support its position that address information is part of an address, OPTi relies on the specification:

In the first address phase, host platform interface circuitry 115 drives out A[23:10] on MAD[15:2] ... In the second address phase, host platform interface circuitry 115 drives out the remaining address information, A[9:0] ... as shown in FIG. 2b.

'141 patent, Col. 6:65-67. Thus, OPTi contends, the specification makes clear that in the disclosed

embodiment two address phases are required to convey a complete address. The first address phase does not completely identify the intended destination or receiver of a transmission as in nVidia's proposed construction because the second address phase contains the remaining address information.

nVidia responds by noting that, while address information may be sent in two portions over two address phases, the '141 claims indicate that the address information identifies the complete intended destination and that anything less than the complete address is only a portion of the address information. nVidia supports its argument by relying on language from claims 4, 8, 15, and 23 from the '141 patent. For example, claim 4 of the '141 patent states:

The computer system of claim 2, wherein said address phase is subdivided into a plurality of sub-phases, wherein each sub-phase carries a portion of said address information.

Thus, nVidia contends, the address identifies the destination and information identifying a part of the address is "a portion of the address information." nVidia asserts that OPTi's proposed construction of address information contradicts claims 4, 8, 15, and 23 of the '141 patent and is therefore presumptively erroneous. *See Southwall Technologies, Inc. v. Cardinal IG Co.*, 54 F.3d 1570, 1579 (Fed.Cir.1995).

nVidia is incorrect. Reading claim 2 in light of claim 4 makes clear that "address information" may be part- or all-of an intended destination. The relevant portion of claim 2 reads as follows:

[A]ddress information is carried on said address-data bus during an address phase of said first cycle and said data information is carried during a data phase of said first cycle ... [.]

Applying nVidia's construction to claim 2 would require that a complete address must fit on the address phase of a first cycle. Thus, nVidia's construction would mandate that claim 2 would *only* apply to addresses that are small enough to fit on a single phase. However, claim 4 makes clear that the address can be broken into a plurality of sub-phases and would permit transfer of addresses of varying size-including addresses larger than 16 bits. *See, e.g.*, '141 patent, Col. 6:7-18. In short, under nVidia's construction, claim 4 would be broader than claim 2. As claim 4 depends from claim 2, this construction would violate the canons of claim construction and is therefore rejected. The Court therefore adopts OPTi's construction for "address information."

4. bus

In a case of role reversal the Court cannot help but notice, the parties have taken atypical approaches to the construction of "bus." OPTi, the plaintiff, has minted a narrow construction of "bus" to mean "a communications interface connecting two or more devices, in which information is conveyed on a set of bi-directional signal lines. A bus is distinct from a point-to-point unidirectional interconnect." nVidia contends that a "bus" is simply "a set of signal lines."

OPTi bases its construction on citations from the preferred embodiment in the specification. In addition, OPTi contends that other claim language requires its "bi-directional" limitation. The Court has previously rejected these arguments when offered by nVidia and rejects them again. OPTi does not deny that nVidia's construction is the well-known construction of one of ordinary skill in the art. Further, OPTi's arguments do not overcome the presumption that a patentee is not limited to the preferred embodiment absent a clear disavowal of claim scope. Accordingly, "bus" means "a set of signal lines."

5. multiplexed/multiplexed manner

Each parties' construction for multiplexed is long and contains numerous limitations. OPTi's construction for "multiplexed" is "individual signal lines carry more than one type of specific information in a time

interleaved manner (i.e., one specified type of information is carried at a first time and a second specified type is carried at a second time). In addition, each specified information type must be carried on a signal line that also carries at least one other specified type of information." nVidia's construction for "multiplexed" is "carrying more than one type of information in a time-interleaved manner using phase techniques, such that address information is carried on the bus during an address phase, data information is carried on the bus during a data phase, and other information, such as cycle definition information, command information, or DMA information, is also carried on the bus during either the data or address phase."

While the parties' constructions appear vastly different, the true "rub" in the alternative constructions is whether "multiplexed" is limited to using phase techniques as proposed by nVidia. Phasing is the method of multiplexing used in the preferred embodiment. However, there are numerous methods of multiplexing—a fact which nVidia does not contest. For the reasons set forth above in Section III.2 regarding the construction of address-data bus, the Court rejects nVidia's attempt to limit multiplexing to phasing techniques. Accordingly, the Court adopts OPTi's proposed construction.

6. *command information*

OPTi contends that "command information" is "information relating to one or more commands for controlling a bus transfer, rather than to address or data." nVidia contends that the proper construction for "command information" is "bits that define one or more commands for controlling the bus transfer, including the bits defining M/IO# , W/R# , and SBHE# rather than address or data bits."

nVidia again seeks to limit the construction of a term to the preferred embodiment. This methodology has been rejected by the Court. While it is clear that M/IO# , W/R# , and SBHE# are commands used in the preferred embodiment, the patentee makes clear that these commands are merely examples of commands and not the universe of commands that can be used. *See* '141 patent, Col. 8:44-45 "command information (e.g., M/IO# , W/R# , SBHE#)" (emphasis added); *see also* Col. 6:64-65 (same). Accordingly, the Court rejects nVidia's construction and adopts OPTi's proposed construction for "command information."

7. *ISA command information*

OPTi offers a plain language construction for "ISA command information." OPTi asserts that "ISA command information" is "information correlating to one or more of the commands recognized in the Industry Standard Architecture bus standard." nVidia asserts that "ISA command information" means "the bits that identify at least SBHE# and one of MEMER# , MEMW# , IOR# , and IOW# signals and may also identify any other command signals as defined by the ISA specification."

As nVidia's construction improperly attempts to incorporate limitations from the preferred embodiment, it is rejected. There is nothing in the specification that would limit the construction of ISA commands in the manner nVidia suggests. nVidia implicitly admits this fact as it principally relies on statements made during prosecution to support its construction. However, a review of the prosecution history makes clear that the commands nVidia seeks incorporated into the construction of "ISA commands" are merely exemplars of ISA commands.

In contrast, the device of the present invention transmits commands that correlate to ISA commands *such as* those shown on pages 2-4 of the present application.... As shown in Table 2 of the present application, the low three bits of the MAD bus can convey information that correlates to *at least the following* ISA signals: SBHE# , MEMR# , MEMW# , IOR# , IOW# .

September 4, 1998 Preliminary Amendment for Continued Prosecution Application at 15 (emphasis added).

In direct contrast to nVidia's proposed construction, OPTi proposes a plain language construction consistent with the specification and the claims. The Court adopts OPTi's proposed construction. "ISA command information" means "information correlating to one or more of the commands recognized in the Industry Standard Architecture bus standard."

B. Pre-Snoop Patents

1. bus master

OPTi's construction of "bus master" is "the device which initiates a data transfer on a bus." nVidia contends the "bus master" is "a device connected to a bus which arbitrates for control of the bus and can receive data transferred on the bus."

Both parties rely on PCI Local Bus Specification Revision 2.0 that was specifically incorporated by reference into the patent specification. In the PCI Local Bus Specification, there are references to the bus master arbitrating for each access it performs. nVidia asserts therefore that its construction must be correct. The Court is not persuaded by nVidia's argument.

The most relevant section of the PCI Local Bus Specification 2.0 to determine a proper construction for "bus master" would be the glossary. A review of the glossary section of the PCI Local Bus Specification 2.0 reveals that the "master" is defined as the "agent that initiates a bus transaction." Further, the specification is consistent with this definition. *See* '906 patent, Col. 10:34-35 ("These data transfers may be either read or write data transfers; the master is the initiator, and the target is the responding device ..."). Accordingly, the Court adopts OPTi's construction of "bus master."

2. sequentially transferring data units between said bus master and said secondary memory

OPTi contends that this means "moving the data units from one place to another; in this instance, between the bus master and the secondary memory, in the sequence in which they are stored." nVidia's construction of the phrase is "moving data units on the bus in sequential order where the transfer of the next data unit from secondary memory begins after the transfer of the previous data unit to the bus master is complete."

As a preliminary matter, the Court notes that the question of what constitutes a data transfer does not appear to be in dispute. OPTi characterizes a "transfer" as "moving data units from one place to another, in this instance, between the bus master and the secondary memory." Likewise, nVidia describes a "transfer" as the process of "moving data units on the bus ... from secondary memory ... to the bus master." These definitions are consistent with one another and with the use of the term in the Pre-Snoop patents.

The primary dispute between the parties appears to be whether the language calling for a sequential transfer means that the transfer of one data unit cannot begin until the transfer of the prior unit is complete as nVidia contends. The principal application for the invention described in the Pre-Snoop patents was "burst" transfers. Sequential ordering in which data units are transferred is what makes the basic mechanism of the burst transfer possible:

The PCI bus achieves very high performance, in part because its basic data transfer mode is by burst. That is, data is always transferred to or from a PCI device in a known sequence of data units defined by a known sequence of data unit addresses in an address space. In the "cache line" burst mode, exactly four transfers take place. In the "linear" burst mode, any number of transfers (including 1) can take place to/from linearly sequential addresses until either the initiator or the target terminated the transaction. In either mode, the initiator need only specify the starting address because both parties know the sequence of addresses which follow.

'906 patent, Col. 5:5-16. This would tend to undermine nVidia's proposed construction. In fact, the Court

can find nothing in the specification or prosecution history of the Pre-Snoop Patents that would limit the construction in the manner nVidia proposes.

nVidia's principal basis for its proposed construction is the assertion that claim limitations must be construed from the perspective of the PCI bus or bus master. The Court finds this argument unpersuasive, albeit novel. nVidia has not pointed to, nor has the Court found, any support in the specification for the requirement that certain limitations be read from the perspective of the bus master or PCI bus. Nor has nVidia cited to any case law to support this position. nVidia's argument is simply an attempt to limit the construction to the preferred embodiment. Such a methodology must be rejected. Accordingly, "sequentially transferring data units between said bus master and said secondary memory" is construed to mean "moving data units from the bus master and the secondary memory in the sequence in which they are stored."

3. next-line

OPTi's construction for "next-line" is "the line immediately following the line being transferred." nVidia proposes that "next-line" be construed as "the line immediately following the line of secondary memory being received by the bus master."

The only dispute between the parties regarding the term "next-line" is whether, during a transfer, the memory can be in transit to the bus master or whether it has already arrived at the bus master. nVidia asserts that a line is the "next-line" only when it is "the line ... immediately following the line ... being received by the bus master" instead of simply being the next line in sequence. For example, if line 1 were in transit from the target to the master, the "next line" to be transferred-i.e., the next line in sequence-and, the next line to be snooped, would be line 2. This is true regardless of whether any of the data units of line 1 had arrived at their destination or not.

nVidia's proposed construction for this term is based on a construction of the limitation from the perspective of the bus master or PCI bus. The Court has rejected that approach. Accordingly, "next-line" is construed to mean "the line immediately following the line being transferred."

4. initiating a next-line inquiry

OPTi contends that the phrase "initiating a next-line inquiry" means "sending a command to snoop the next line." nVidia asserts that "initiating a next-line inquiry" means "sending a command to snoop the next-line (e.g., asserting EADS #), without an associated request to transfer the corresponding line from secondary memory."

The parties again offer very similar constructions. The only issue for the Court to decide for this construction is whether "initiating a next-line inquiry" requires the additional limitation nVidia offers of "without an associated request to transfer the corresponding line from secondary memory."

There is no basis in the language of the claims to support nVidia's construction. Nor has nVidia pointed to a clear disavowal in the specification or prosecution history that would require adoption of this negative limitation. Accordingly, it is rejected. The Court adopts OPTi's proposed construction.

5. initiating a next-line inquiry ... to determine whether an N+Ith byte line of said secondary memory is cached in a modified state in said first cache memory

OPTi proposes that this phrase means "initiating a next-line inquiry to determine whether the next-line has been modified by the CPU." nVidia responds by asserting that the proposed construction of the phrase should be "initiating a next-line inquiry to determine whether the next-line has been modified by the CPU or may be read by the bus master without any further concern that the cache contains a more current copy of the data."

Once again, the parties have narrowed the dispute to a single question. In this case, the Court need only decide whether the additional limitation "or may be read by the bus master without any further concern that the cache contains a more current copy of the data" is required.

nVidia contends that OPTi's construction could allow for stale data as it is not limited in time. Thus, nVidia seeks to add the additional limitation recited above. nVidia's concern has merit as the goal of the invention is to prevent the transfer of stale data. However, nVidia's proposed construction would likely confuse the jury and is therefore rejected. The Court instead construes "initiating a next-line inquiry ... to determine whether an N+1th byte line of said secondary memory is cached in a modified state in said first cache memory" to mean "initiating a next-line inquiry to determine whether the next-line of data in the secondary memory is different from the corresponding data in the cache."

6. concurrently with at least one of the data unit transfers in said step of sequentially transferring

The parties generally agree on the proper construction of this phrase. OPTi contends that the construction should be "while at least one of the data units is moving from the secondary memory to the bus master." nVidia's construction is similar: "while at least one of the data units is moving directly from the secondary memory to the bus master." Thus, the only dispute between the parties for the aforementioned phrase is whether the movement from the secondary memory to the bus master is limited to a direct transfer.

There is no requirement in the claims that the transfer be direct. Further, what support nVidia relies on in the specification relates to the preferred embodiment. The Court will not read limitations of the preferred embodiment into the claims. Accordingly, the Court adopts OPTi's construction.

7. means for sequentially transferring data units between said bus master and said secondary memory beginning at a starting memory location address in said secondary memory address space ... said sequentially transferred data units including a last data unit before said I-byte boundary and a first data unit beyond said I-byte boundary

The parties agree that the aforementioned phrase is subject to 35 U.S.C. s. 112, para. 6. Pursuant to that statutory section, the parties have also agreed on the correct function. Thus, the only issue the Court must decide is the structure that performs the function.

OPTi's original proposed structure was "a PCI bus operating in linear incrementing burst mode under the control of a system controller." nVidia's proposed structure is "the system controller (SYSC) and integrated peripherals controller (IPC) (able to recognize FRAME# , IRDY# and TRDY# signals and operate pursuant to Figs. 4-7), PCI bus, and control lines between the SYSC/IPC and secondary memory." OPTi has now amended its construction to be "the system controller (SYSC) and integrated peripherals controller (IPC), a PCI bus, and control and data lines between the SYSC/IPC and secondary memory." In short, OPTi now proposes a construction very similar to that of nVidia and it is not clear whether there is even a real dispute between the parties. Assuming the parties are not in agreement, the Court address the structure below.

The corresponding structure for the recited function must include all of the structure disclosed in the specification that is necessary to perform the recited function. *See* *Cardiac Pacemakers, Inc. v. St. Jude Medical, Inc.*, 296 F.3d 1106, 1119 (Fed.Cir.2002). Figure 1 of the '906 patent, shown below, reveals what structure is necessary to sequentially transfer data units between the bus master and secondary memory. As can be seen from the Figure, the structure necessary to transfer data between secondary memory and the bus master conforms to the structure proposed by OPTi.

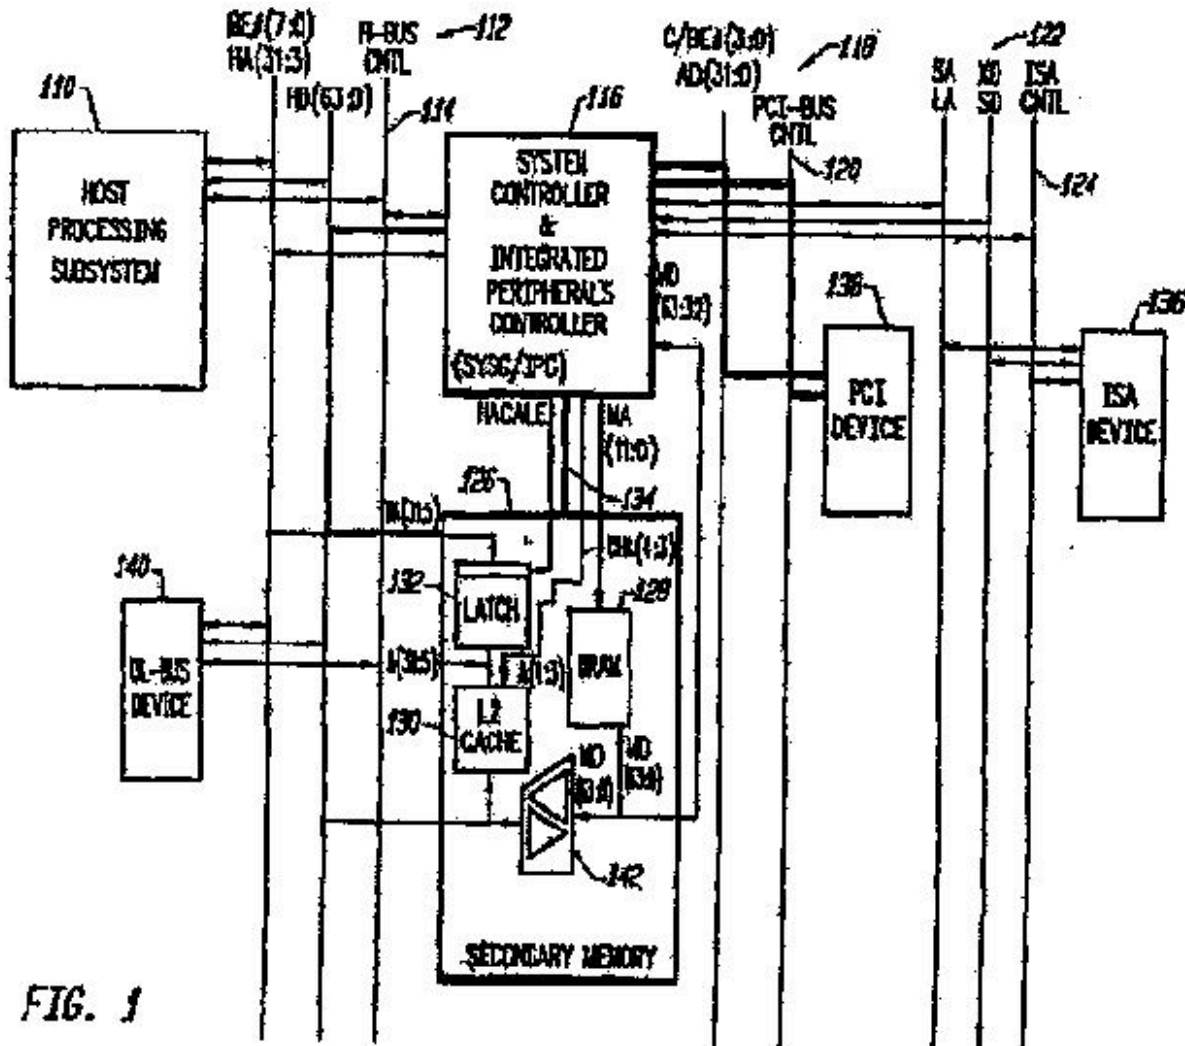


FIG. 1

Accordingly, the structure linked to the means for sequentially transferring data limitation is "the system controller (SYSC) and integrated peripherals controller (IPC), a PCI bus, and control and data lines between the SYSC/IPC and secondary memory."

8. means for initiating a next-line inquiry, prior to completion of the transfer of the last data unit before said I-byte boundary, to determine whether an N+I'th I-byte line of said secondary memory is cached in a modified state in said first cache memory, said N+I'th I-byte line being a line of said secondary memory which includes said first data unit beyond said I-byte boundary

The parties agree here, as they did above, that the recited phrase is written in means-plus-function language and is subject to 35 U.S.C. s. 112, para. 6. Again, the parties also agree on the function. However, the parties' positions regarding the corresponding structure are very different.

OPTi contends that the corresponding structure is "the logic circuitry of the SYSC/IPC schematically illustrated in Figure 9 of the patent." Originally, nVidia proposed that the corresponding structure was "SYSC/IPC and host processing subsystem. The SYSC/IPC must be capable of generating the following signals: (1) HOLD, HLDA, AHOLD, or BOFF# , (2) inquire address, (3) INV, and (4) EADS# , and operate

pursuant to Figs. 4-7. HOLD, HLDA, and AHOLD are generated by the circuitry in Fig. 11. EADS# is generated by the structures in Figs. 8, 9." nVidia now contends that the patent is invalid for failure to disclose any corresponding structure.

nVidia's new position is one this Court has seen before. It is nVidia's position that the corresponding structure for performing the function is software. "The Federal Circuit has made clear that when software is linked to the disclosed function, the structure for performing that function is linked to the algorithm disclosed in the specification." *See Gobeli Res., Ltd. v. Apple Computer, Inc.*, 384 F.Supp.2d 1016, 1022 (E.D.Tex.2005).

nVidia asserts that a necessary part of the corresponding structure for the disclosed function is the structure associated with incrementing the snoop address. However, the patent specification makes clear that this information is not disclosed in the Pre-Snoop patents:

The output of NAND gate 910, FTRDTGB, is connected to the D input of a flip-flop 912, which is clocked on LCLK1. Flip-flop 912 thus delays FTRDTGB by one PCICLK to enable other circuitry (not shown) in the system controller 116 to increment the secondary memory line address on HA(31:5) (FIG.1) ...

'906 patent, Col. 24:10-15. nVidia contends that because the structure for incrementing the address for the next-line inquiry is "not shown," the algorithm cannot be described. Accordingly, nVidia contends, any claims of the Pre-Snoop patent containing the aforementioned phrase are invalid as a matter of law. *See Gobeli*, 384 F.Supp.2d. at 1023.

nVidia is incorrect. nVidia is asserting structure linked to a function of actual implementing a snoop rather than the recited function of initiating a snoop. The patent makes clear that the "next-line" inquiry is "initiated" by the PSNSTR1 signal. '906 patent Col. 23:34-36 ("PSNSTR1 carries a high-going pulse when it is desired to initiate a predictive snoop cycle during a PCI master burst transfer"); Col. 24:32-35 ("As previously described, PSNSTR1 is provided to an input of NAND gate 822 in FIG. 8 and, like LT2, initiates an L1 cache inquiry cycle"). Further, the patent identifies the structure that generates PSNSTR1:

FIG. 9 is a schematic diagram of circuitry in the system controller 116 which produces the PSNSTR1 signal used in FIG. 8. As previously mentioned, PSNSTR1 carries a high-going pulse when it is desired to initiate a predictive snoop cycle during a PCI master burst transfer.

'906 patent, Col. 23:32-36; *see also id.*, Col. 23:37-24:35 and Fig. 9. Accordingly, the Court finds that the Pre-Snoop patents do disclose corresponding structure for the recited function and adopts OPTi's proposed construction.

E.D.Tex.,2006.
OPTi Inc. v. nVidia Corp.

Produced by Sans Paper, LLC.